



A Robust Controlled Backward Reach Tube with (Almost) Analytic Solution for Two Dubins Cars

Ian M. Mitchell *

Department of Computer Science
The University of British Columbia
Vancouver, British Columbia, Canada
ian.mitchell@ubc.ca

Abstract

Benchmark Proposal: We describe how a well-known backward reachability problem with nonlinear dynamics and adversarial inputs—based on a pursuit evasion game with two identical vehicles that have Dubins car dynamics—can be viewed as a robust controlled backward reach tube. The resulting set is nonconvex with a surface that is nondifferentiable in places, yet (mostly explicit) closed form solutions for points on the surface of this set have been derived based on a classical differential game analysis, and so these points can be sampled with high accuracy at arbitrary density. We propose this problem as a benchmark because few existing reachability algorithms can tackle robust controlled backward reach tubes despite their potential for proving the robust safety of systems, and this (almost) analytic solution exists against which to compare prospective solutions. We then describe some extensions to the problem to provide additional future challenges. Code is provided.

1 Introduction

The controlled invariant set is a common tool used to prove that systems can be controlled to remain within some set of states, and the robust controlled invariant set is a version in which the invariance can be shown despite perturbation of the dynamics. Similar problems have been studied in reachability, although differentiating between research on reachability algorithms which work for systems with no inputs, with a single input, or with two adversarial inputs has been a challenging task because all use the generic term “reachability”. In [Mitchell, 2011] we proposed notation to differentiate these cases—as well as whether the set includes states reachable over an interval of time or at a single time, and whether the computation proceeds forward in time from an initial set or backward in time from a target set—but the notation is cumbersome, and unnecessary when considering algorithms which focus on solving a single version of the problem. More importantly, notation cannot easily appear in titles, abstracts or introductions, making it hard for others to easily identify which particular problem is being solved in a given paper.

*Research supported by the [National Science and Engineering Research Council of Canada](#) (NSERC) Discovery Grant #4543.

Increasing use of the word “tube” for the case in which reachability is evaluated over an interval of time has helped the situation, so to differentiate between reachability problems with various treatments of inputs we propose to follow the lead of robust controlled invariant sets and use the phrase “robust controlled” to identify the category of reachability problems with adversarial inputs. In particular, in this paper we propose a benchmark problem for *robust controlled backward reach tubes* (RCBRTs), in which we seek to compute the set of states giving rise to trajectories which can be controlled by the input(s) of one player to reach a given target set at a some time during a given time interval in a manner robust to the disturbance of the second player’s input(s) on the dynamics. RCBRTs clearly share features with robust controlled invariant sets as well as discriminating kernels from viability theory [Aubin et al., 2011], with the key difference being the presence of a terminating target set (similar to viability theory’s capture basins).

Impressive advances have been demonstrated in the scalability and accuracy of reachability algorithms for systems with fewer inputs over the past decade; for example [Bak et al., 2019] demonstrates a previously unimaginable billion dimensional problem, albeit for a system with linear dynamics and no inputs. Much less progress has been made on RCBRTs; for example, [Kaynama et al., 2015] demonstrates an analysis on a twelve dimensional linear system with adversarial inputs.

Unlike reachability analyses without adversarial inputs, RCBRTs do not appear to be amenable to efficient sampling approaches because (as demonstrated below) distinct optimal trajectories may arise from a single point in the state space and/or lead to a single point in the target set; consequently, two-point boundary value based or optimization based sampling may miss such trajectories, and dense sampling of the competing inputs from all possible initial states is generally infeasible.

And yet RCBRTs have considerable potential in analysis and synthesis of safe controllers which will drive a system toward (or away from) target sets despite uncertainty in the model, environment, or actions of other agents. Examples beyond the collision avoidance problem considered here include automated delivery of anesthesia [Yousefi et al., 2019], learning a Gaussian process model of drone flight while maintaining safety [Fisac et al., 2018], and analysis of the user interface of an automated landing system [Oishi et al., 2008].

With this benchmark proposal we hope to drive interest in developing more general and still scalable algorithms for RCBRTs. The contributions of this benchmark proposal are to

- Describe an RCBRT whose boundary can be sampled analytically to arbitrary density. The analytic solution can be used to validate and analyze the accuracy of potential solutions. The RCBRT:
 - Is derived from a collision avoidance problem involving two Dubins cars.
 - Features nonlinear dynamics with bounded, adversarial inputs.
 - Is nonconvex with a surface that is non-differentiable in places.
 - Can be characterized in either three or six dimensions depending on whether a relative coordinate frame is used.
 - Includes pairs of initial set / RCBRT points connected by multiple distinct optimal trajectories.
- Describe extensions to the basic problem for which the system dimension is larger. For some of these extensions, the analytic solution for the basic problem can be used to find under or over approximations of the true RCBRT.

The paper is organized as follows: Section 2 describes the problem for which the (almost) analytic solution is known, and how that solution can be characterized. Section 3 describes a number of extensions to the basic problem and how they relate to the (almost) analytic

solution. Section 4 briefly describes how these problems can be formulated as Hamilton-Jacobi-Isaacs PDEs whose solutions can be numerically approximated, although these techniques scale poorly with dimension. Source code can be found at https://bitbucket.org/ian_mitchell/arch-2020-benchmark.

2 The Basic Version

In this section we describe the problem which leads to an RCBRT in three dimensions and summarize an analysis which allows us to write closed form solutions for points on the boundary of the tube. All of these elements have been examined in prior publications, so the purpose of this section is to provide a cohesive overview and justification for use as a benchmark rather than the details of the constructions; however, those details are available in the citations for interested readers.

2.1 The Problem

The motion of a kinematic unicycle model is governed by the ODE

$$\dot{x} = \frac{d}{dt} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} v \cos x_3 \\ v \sin x_3 \\ \omega \end{bmatrix}$$

where

- The *state* $x \in \mathbb{R}^2 \times [-\pi, \pi)$ consists of the position in the plane (x_1, x_2) and heading x_3 .
- The *inputs* are the linear velocity $v \in [\underline{v}, \bar{v}] \subseteq \mathbb{R}$ and angular velocity $\omega \in [-\bar{\omega}, \bar{\omega}] \subseteq \mathbb{R}$. Note that we assume symmetry of the angular velocity range.

In a Dubins car model [Dubins, 1957] the linear velocity is a fixed constant $v = \underline{v} = \bar{v}$, often chosen as $v = +1$, while angular velocity is allowed to vary. In a Reeds-Shepp car model [Reeds and Shepp, 1990] the linear velocity is chosen from two symmetric fixed values; for example, $v \in \{-1, +1\}$. Both cases can equivalently be defined as a car with limited turn radius (equal to $\bar{\omega}^{-1}$ for unit linear velocity), whereas the unicycle can turn in place if the linear velocity can be set to zero.

Numerous works have examined the reach sets of individual vehicles traveling under various versions of unicycle dynamics—for example, in [Patsko and Turova, 2009] the authors examine the case where $\bar{v} = +1$ and $\underline{v}$ is allowed to vary—but in these cases all nondeterminism in the input signal(s) seeks the same outcome (typically to make the reach set larger).

In this paper we consider an adversarial game scenario involving two identical vehicles in which the input(s) of one are trying to bring the vehicles close together while the input(s) of the other are trying to keep them apart. It can be thought of as a version of the homicidal chauffeur problem [Isaacs, 1999], but with each of the two agents having identically limited unicycle dynamics rather than one being a car and the other a pedestrian.

In order to distinguish between the two vehicles we will call the one seeking to drive the vehicles together the *pursuer* and the one seeking to keep them apart the *evader*. In this basic problem we will fix the linear velocity $v = +1$ (so the vehicles behave as Dubins cars) and the bounds on the angular velocity as $\bar{\omega} = +1$. We will declare that the pursuer wins if the vehicles pass within distance $\beta > 0$ of each other in the plane, regardless of their headings; otherwise, the evader wins.




Figure 1: Relative coordinate system with the evader at the origin.

Because the winning conditions depend only on the relative position of the vehicles, we can recast the problem in relative coordinates. For reasons involving the algorithm discussed in section 4, we choose to put the evader at the origin of the relative coordinate system; however, the analytic solution can be derived with either vehicle at the origin (see section 2.2).

In the relative coordinate system shown in figure 1, the dynamics of the system are given by the nonlinear ODE

$$\dot{x} = \frac{d}{dt} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} -v_e + v_p \cos x_3 + \omega_e x_2 \\ v_p \sin x_3 - \omega_e x_1 \\ \omega_p - \omega_e \end{bmatrix}, \quad (1)$$

where

- $(x_1, x_2) \in \mathbb{R}^2$ is the relative position in the plane.
- $x_3 \in [-\pi, \pi)$ is the relative heading,
- v_e and v_p are the linear velocities of the evader and pursuer respectively. These velocities are constant and equal so they are not considered inputs in this version of the problem.
- ω_e and ω_p are the angular velocities of the evader and pursuer respectively. These are drawn from the same interval $[-\bar{\omega}, \bar{\omega}]$ and are the inputs for this version of the problem.

We will define the domain in which we compute the RCBRT as

$$\mathcal{X} = \mathbb{R}^2 \times [-\pi, \pi) \quad (\text{or } \mathbb{R}^2 \times \mathbb{S}^1) \quad (2)$$

using this relative coordinate system.

The backward reach tube that we seek to compute represents the set of states from which the pursuer can drive the system trajectory into the *target set*

$$\mathcal{T} = \left\{ x \in \mathcal{X} \mid \sqrt{x_1^2 + x_2^2} \leq \beta \right\}. \quad (3)$$

In the three dimensional state space, $\mathcal{T}$ is a cylinder because it does not depend on relative heading x_3 . Figure 2 shows $\mathcal{T}$.

We consider the set to be computed an RCBRT because the pursuer will use its input ω_p to try to control the system state toward $\mathcal{T}$, but must do so in a way which is robust to the efforts of the evader to steer away through its input ω_e . Because we are operating in continuous time, the question arises of what each vehicle knows about the other's choices. We




Figure 2: Two views of the surface of the RCBRT (red shape) for the basic problem. On the left, the RCBRT surface is partially transparent and the target set is shown as a solid blue cylinder. This set is the same as that shown in [Mitchell et al., 2005, Figure 3], [Mitchell, 2002, Figure 3.5] and [Mitchell, 2001, Figure 8], but with relative heading in $[-\pi, \pi)$ (instead of $[0, 2\pi)$) and different values for β , v_e and v_p .

will give any potential informational benefit to the pursuer by allowing it to use a nonanticipative strategy. The intuition for how this choice affects the problem is that at a given time instant the evader knows the current state of the system and how the pursuer will respond to any given evader input, while the pursuer knows both the state and the evader's actual chosen input. The notation for non-anticipative strategies becomes somewhat cumbersome, so the interested reader is referred to [Mitchell et al., 2005, section 2] for more details.

Remark 1. The dynamics (1) are separable, meaning they can be written in the form

$$\dot{x} = f_0(x) + f_e(x, \omega_e) + f_p(x, \omega_p),$$

where each of the component terms depends on no more than one of the inputs ω_e and ω_p . While the details of the non-anticipative strategy information pattern is important to achieve mathematical rigour for cases with general dynamics, for separable dynamics it is equivalent to the more intuitive state feedback information pattern; in other words, each player's instantaneous optimal input does not depend on the other player's instantaneous choice.

The RCBRT is formally defined as

$$\mathcal{R}([-T, 0], \mathcal{T}) = \{x \in \mathcal{X} \mid \exists \omega_p(\cdot), \forall \omega_e(\cdot), \exists s \in [-T, 0], x(s) \in \mathcal{T}\} \quad (4)$$

where

- The (backward) horizon is $T > 0$.
- An input signal $u(\cdot) : [T, 0] \rightarrow \mathcal{U}$ is drawn from the set of measurable functions mapping a time to a valid input value.
- With an abuse of notation, the pursuer's nonanticipative strategy $\omega_p(\cdot)$ is a mapping from evader input signals to pursuer input signals.
- The evader's input signal $\omega_e(\cdot)$ is just a regular input signal.

The definition (4) may appear to give the information benefit to the evader since $\omega_e(\cdot)$ is picked last, while at the same time restricting both input signals to be open loop since they are both functions mapping time to input value. However

- Because $\omega_p(\cdot)$ is a strategy it can respond to any particular choice of $\omega_e(t)$ instantaneously even though $\omega_e(\cdot)$ is chosen later.
- Because $\omega_p(\cdot)$ is chosen first, both players know how the pursuer will respond to any given choice by the evader; hence both players can deduce the current state at any given time and thereby have access to the necessary information to build closed loop input signals.

Measurable functions and non-anticipative strategies may seem like mathematically heavy artillery when we might instead desire to impose more practical constraints on the input signals for real problems, such as piecewise constant or linear state feedback. In fact, the optimal input signals needed for the analytic solution of this basic problem turn out to be piecewise constant (see section 2.2), but we call out the big guns here in order to enable the approximation of this RCBRT and the RCBRTs for extended versions of this problem through viscosity solutions of HJI PDEs in section 4.

At the same time, the need for these precise definitions should serve as a warning that the definition and solution of this RCBRT is mathematically delicate in places. In particular, the target set $\mathcal{T}$ is the closure of an open set, and in this case it can be shown that the RCBRT is also a closed set (for example, see [Mitchell et al., 2005]). Given that the definitions all have obvious complements, one might be tempted to deduce that the evader's RCBRT is precisely the complement of the pursuer's RCBRT (4); however, the story is more complicated,¹ and the interested reader can find a rigorous treatment in [Cardaliaguet et al., 1999, section 4]. Fortunately, those interested in computing over and under approximations numerically can use the benchmark formulation above and for the most part avoid these mathematical technicalities.

2.2 An (Almost) Analytic Solution

In [Merz, 1972] the author deduces the optimal trajectories for the problem described above, albeit for the relative coordinate system with the pursuer at the origin. In [Mitchell, 2001] we write out Merz's equations in detail for his case, and then replicate the analysis for the relative coordinate system with the evader at the origin. We provide accompanying MATLAB code which regenerates the figures and can compute points on the boundary of the RCBRT to an accuracy comparable with the available implementation of the basic trigonometric functions.

Given the detailed exposition available in these manuscripts, the purpose of this section is merely to highlight interesting characteristics and features of this analytic solution, and the interested reader can find the details in [Merz, 1972, Mitchell, 2001].

In the original paper, Merz parameterizes trajectory terminal points on the boundary of the cylindrical target set $\mathcal{T}$, and then claims that points on the *barrier* (the portion of the boundary of the RCBRT which is not also the boundary of the target set) can be described as the initial conditions of one or more optimal trajectories made up of segments along which both inputs are fixed. Three types of parameterized segment exist, each having two versions depending on the choice of input:

- Opposite inputs: In this type $\omega_e = -\omega_p$, and the two versions depend on whether $\omega_e = +1$ or $\omega_e = -1$.
- Same inputs: In this type $\omega_e = \omega_p$, and the two versions depend on whether $\omega_e = +1$ or $\omega_e = -1$.
- Pursuer not turning: In this type $\omega_p = 0$, and the two versions depend on whether $\omega_e = +1$ or $\omega_e = -1$.

¹In particular, [Mitchell et al., 2005, Lemma 8, case 2] has a circular dependency which breaks the proof (much thanks to Andre Platzer for spotting this error), but the version of the same lemma in [Mitchell, 2002, section 2.4.3] works fine with a slightly weakened claim containing a strict inequality.




Figure 3: Two views of slices at fixed x_3 of the boundary of the RCBRT with $\beta = 0.5$ and the evader at the origin. Note that the x_1 and x_2 coordinates are swapped compared to the coordinate system in figure 1. The large circles are slices of the target set $\mathcal{T}$. The star symbols are the crossover points for each slice, the small circle symbols are the intersection of the barrier and the target set slices, and the triangle / square symbols denote points on the barrier slices where the combination of optimal trajectory segments to reach the barrier changes. Figure taken from [Mitchell, 2001, figure 6].

The boundary of the RCBRT can be sliced along fixed values of x_3 (horizontal slices in figure 2), and the boundary of each slice divided into three smooth curves: Some half of the circle which is the corresponding slice of the target set $\mathcal{T}$, and two curves which start off tangentially from the two ends of the half circle and meet at a *crossover point* where the barrier is not differentiable.

The first step to compute points lying on a particular slice of the barrier is to determine the crossover point. There are two possible optimal trajectories arising from this point on the left and two on the right depending on which trajectory segments are involved, and the crossover point can be found by equating each possible pair of left and right, and solving the resulting two equations (one each for coordinates x_1 and x_2 since the slice x_3 is fixed) for two unknowns (the time durations of the two trajectories) to find the combination which is physically realizable.

The remainder of the points on the barrier arise from one of eight possible combinations of trajectory segments, with a left-right asymmetry arising from the complementary shapes of the RCBRT above and below $x_3 = 0$. Each combination has a free parameter, so by sampling these free parameters and the slices in x_3 it is possible to generate points on the barrier to an arbitrary density.




Figure 4: Two views of slices at fixed x_3 of the boundary of the RCBRT with $\beta = 0.5$ and the pursuer at the origin. The labelling is the same as that in figure 3. Figure taken from [Mitchell, 2001, figure 5], and the left subplot recreates [Merz, 1972, figure 7].

Explicit, closed form equations for all of the points on the barrier can be written except for the crossover points. The crossover points have closed form equations, but they must be solved numerically because they are not explicit and involve nonlinear trigonometric terms; consequently, we call Merz’s solution “almost analytic.” In practice the equations can be solved to arbitrary accuracy using appropriate numerical methods, and are solved to nearly IEEE double precision floating point accuracy in the accompanying MATLAB code.

Figure 3 shows two views of slices of the RCBRT at

$$x_3 = -\left(\frac{k}{6}\right)\pi \quad \text{for } k = \{1, 2, 3, 4, 5, 6\}$$

with the evader at the origin and for $\beta = 0.5$. Note that the x_1 and x_2 axes are swapped for figures in this section compared to the axes defined in figure 1 (the swapped axes agree with the coordinate system from [Merz, 1972]). The x_3 axis is also labelled in degrees rather than radians. The figure shows that the RCBRT is not smooth at the crossover points. It is also possible to deduce that the RCBRT is nonconvex: Consider the line joining $x = (0, 0.5, 0)$ on the top slice of the target set to the point $x = (0, 2.9, -\pi)$ just below the crossover point for the bottom slice. This line will pass through the point $x = (0, 0.9, -\pi/6)$, which is clearly not within the RCBRT slice for $x_3 = -\pi/6$.

For comparison, figure 4 shows two views of the same slices of the RCBRT with the pursuer at the origin.

Figure 5 shows a coarsely sampled set of optimal trajectories leading to one side of the (pursuer at the origin) barrier for $x_3 = -5\pi/6$. The set of trajectories closer to the origin terminate at different points on the target set and comprise of only a single optimal trajectory




Figure 5: Two views of a coarse sampling of trajectories whose endpoints form a slice of the left barrier of the RCBRT with $\beta = 0.5$ and the pursuer at the origin. Note that the set of trajectories starting from points closer to the crossover point all share a terminal trajectory segment. Figure taken from [Mitchell, 2001, figure 4].

segment. The set of trajectories closer to the crossover point all comprise of two optimal segments, the second of which is shared and terminates at a single point on the target set. Finally, no portion of this side of the barrier for this slice is itself an optimal trajectory segment.

Figure 6 shows the two trajectories leading from the crossover points for two distinct slices (again for the case with the pursuer at the origin). Notably, the trajectories leading from the crossover point for the slice at $x_3 = -5\pi/6$ not only diverge at first, but they then converge to the same terminal point on the target set. The slice at $x_3 = -\pi/2$ shows distinct features. First, the trajectories leading from the crossover point terminate at different points on the target set. Second, a portion of the left barrier is composed of an optimal trajectory segment.

Remark 2. From the problem definition in section 2.1 and the (almost) analytic solution described here we can identify a number of features which make this reachability problem challenging:

- The dynamics (1) are nonlinear, with x_1 and x_2 coupled to all dimensions except themselves.
- The domain (2) is not of the form $\mathbb{R}^d$ for some d because x_3 is periodic.
- The cylindrical target (3) couples x_1 and x_2 , although it could be over or under approximated by a square prism if decoupling is desired.
- There are two inputs with competing goals. The dynamics are separable in the inputs (see remark 1), so we can avoid some complications around non-anticipative strategies; however, they are affine in the inputs and so discontinuous (or *bang-bang*) controls are optimal (see section 4).
- The RCBRT is nonconvex and asymmetric, although there is an odd symmetry across the $x_3 = 0$ plane.




Figure 6: Left and right trajectories leading from crossover points for slices $x_3 = -5\pi/6$ and $x_3 = -\pi/2$ of the RCBRT with $\beta = 0.5$ and the pursuer at the origin. Slices of the target cylinder are also shown at the terminal points of the trajectories. For the slice at $x_3 = -\pi/2$, the portion of the left barrier between the star symbol at the crossover point and the square symbol is part of the left trajectory. Figure taken from [Mitchell, 2001, figure 3].

- The surface of the RCBRT is mostly smooth, but is non-differentiable along the curve of the crossover points.
- Optimal trajectories which lead from points on the boundary of the RCBRT to points on the boundary of the target set arise from piecewise constant input signals with a bounded number of switches for any single trajectory; therefore, it should be possible to compute the RCBRT using only this class of input signals even though it is defined (4) using nonanticipative strategies over measurable input signals.
- Multiple optimal trajectories can start from a single point on the boundary of the RCBRT, terminate at a single point on the target set, or both. Multiple optimal trajectories may share segments. These features are likely to cause problems for algorithms which use two point boundary value solvers to search for optimal trajectories.

We hope that readers will look upon these features as hurdles to overcome through design of novel algorithms rather than insurmountable obstacles.

3 Problem Extensions

In this section we sketch out a few extensions to the basic problem, mostly with the goal of increasing the dimension of the space in which the RCBRT must be found. In some cases an analytic solution based on the solution of the basic problem can be deduced, while in others that solution only provides some guidance.

It is also possible to break the assumption that the vehicles are identical by giving them different bounds on their linear and/or angular velocity inputs. While potentially valuable for solving real world pursuit evasion games, we do not consider this option further in the benchmark because it breaks the derivation of the analytic solution.

3.1 The No-Heading Projection

The “extension” in this section is actually a dimensionality reduction. As part of a general exploration of how RCBRTs can be approximated by their projections into lower dimensional subspaces of the state space, in [Mitchell and Tomlin, 2003, section 3.5] we consider the projection of the basic problem into the x_1 - x_2 subspace; in other words, the two vehicles know their relative position, but not their heading. Since $x_3 \in [-\pi, \pi)$, we can model x_3 as an input to the two dimensional system in the x_1 - x_2 subspace.

If we treat x_3 as an input that the evader can choose, then we are seeking the set of states from which the evader can be captured for any possible relative heading. For identical vehicles the evader can choose relative heading $x_3 = 0$ and always escape unless the system starts inside the target set; in other words, the two dimensional RCBRT is the target set. So this case is not interesting as a benchmark problem.

If we treat x_3 as an input that the pursuer can choose, then we are seeking the set of states from which the evader might be captured for some possible relative heading. The optimal choice of x_3 will depend on the relative position of the pursuer, but it is easy to see that the two dimensional RCBRT will be the projection along x_3 of the basic problem’s RCBRT. Consequently, an analytic solution for this *pursuer’s choice* RCBRT can be found by projection of the solution described in section 2.2.

This simplification can be used in concert with the extensions in the remainder of this section to provide RCBRT problems in a range of dimensions. We will note in passing that problem versions which project out either x_1 or x_2 yield RCBRTs which are either trivial or unbounded, so these versions do not appear to be interesting as benchmark problems either.

3.2 Absolute Coordinates

We used relative coordinates for the two vehicles in section 2 because we could reduce the dimension of the state space from six to three. But if we want a higher dimensional version of the benchmark, we can work directly in the absolute coordinates of the two vehicles. Letting x_1, x_2 and x_3 be the coordinates of the evader and x_4, x_5 and x_6 the coordinates of the pursuer, we can substitute

$$\begin{aligned}\mathcal{X} &= \mathbb{R}^2 \times [-\pi, \pi) \times \mathbb{R}^2 \times [-\pi, \pi) \\ \mathcal{T} &= \left\{ x \in \mathcal{X} \mid \sqrt{(x_3 - x_1)^2 + (x_4 - x_2)^2} \leq \beta \right\}.\end{aligned}$$

for (2) and (3) respectively, and

$$\dot{x} = \frac{d}{dt} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \end{bmatrix} = \begin{bmatrix} v \cos x_3 \\ v \sin x_3 \\ \omega_e \\ v \cos x_6 \\ v \sin x_6 \\ \omega_p \end{bmatrix} \quad (5)$$

for (1) to create a benchmark problem in six state dimensions. Analytic slices of the RCBRT for any fixed evader state can be evaluated by transforming into relative coordinates and using the results from section 2.2. Furthermore, a four dimensional problem can be defined by using absolute coordinates with the no-heading projection from section 3.1.

3.3 More Vehicles

An easy way to increase the dimension of this pursuit-evasion game is to introduce more vehicles. Writing down the dynamics is a straightforward replication of (1) or (5) depending on whether relative or absolute coordinates are used. The definition of RCBRTs (4) requires that each vehicle be assigned to either the evader or pursuer team. Most importantly, some care must be taken to create a target set whose RCBRT will have appropriate semantics.

Single evader. This case is simplest to formulate in either relative coordinates (with the evader at the origin) or absolute coordinates. The target set is the union (if any pursuer can capture) or intersection (if all pursuers must capture simultaneously) of states where a pursuer is sufficiently close to the evader; for example, to encode the case where either of two pursuers can capture in relative coordinates, use

$$\mathcal{T} = \left\{ x \in \mathcal{X} \mid \min \left(\sqrt{x_1^2 + x_2^2}, \sqrt{x_4^2 + x_5^2} \right) \leq \beta \right\}, \quad (6)$$

where (x_1, x_2, x_3) are the relative coordinates of the first pursuer and (x_4, x_5, x_6) are the relative coordinates of the second pursuer. Encoding other semantics requires additional capabilities; for example:

- State constraints (aka *avoid* sets) allow problems in which the pursuers cannot get too close together, as well as problems in which there are obstacles in the environment.
- A hybrid state model can capture problems in which more than one pursuer must capture but they can do so sequentially rather than all at once.

Single pursuer. This case is also easy to formulate in either relative coordinates (with the pursuer at the origin) or absolute coordinates. Requiring a single pursuer to capture multiple evaders simultaneously seems overly challenging, but a target allowing the pursuer to capture any single evader can easily be encoded; for example, using (6) for the situation where there are two evaders and the state variables are the coordinates of those evaders relative to the pursuer at the origin. If the evaders must avoid one another, and we are willing to count an evader collision as a pursuer win, then the target can be defined as, for example,

$$\mathcal{T} = \left\{ x \in \mathcal{X} \mid \min \left(\sqrt{x_1^2 + x_2^2}, \sqrt{x_4^2 + x_5^2}, \sqrt{(x_1 - x_4)^2 + (x_2 - x_5)^2} \right) \leq \beta \right\},$$

where the third term in the minimum encodes that the evaders have gotten too close together. A hybrid state model can be deployed if the pursuer must capture (a subset of) evaders in some sequence.

Multiple pursuers and evaders. Absolute coordinates make more sense in this case because there is no unique vehicle. Encoding an “any pursuer captures any evader” winning condition is straightforward, but an “every evader is simultaneously captured” winning condition can be managed; for example, with two pursuers and two evaders

$$\mathcal{T} = \left\{ x \in \mathcal{X} \mid \max \left[\begin{array}{l} \min \left(\sqrt{(x_1 - x_7)^2 + (x_2 - x_8)^2}, \sqrt{(x_1 - x_{10})^2 + (x_2 - x_{11})^2} \right), \\ \min \left(\sqrt{(x_4 - x_7)^2 + (x_5 - x_8)^2}, \sqrt{(x_4 - x_{10})^2 + (x_5 - x_{11})^2} \right) \end{array} \right] \leq \beta \right\},$$

where

- (x_1, x_2, x_3) and (x_4, x_5, x_6) are the two evaders' absolute coordinates.
- (x_7, x_8, x_9) and (x_{10}, x_{11}, x_{12}) are the two pursuers' absolute coordinates.
- The first minimum encodes the fact that the first evader is captured by one of the pursuers, while the second minimum encodes the fact that the second evader is captured by one of the pursuers. Note that one pursuer can capture both evaders.
- The maximum encodes the fact that both evaders must be captured simultaneously.

Enumeration of all winning simultaneous combinations for the pursuers, including cases where evader collision causes a pursuer win, can be encoded similarly, while winning sequential combinations can be encoded with a hybrid state space.

Analytic solution. The higher dimension and the interactions among the vehicles makes it difficult to imagine how Merz's approach could be extended to provide analytic solutions to these cases; however, we can under or perhaps over approximate some of these solutions using the solution from section 2.2 and pairwise consideration of the pursuers and evaders. For example, if "capture any evader" is a winning condition for the pursuer(s), then every pursuer-evader pair whose relative coordinates fall within the RCBRT for the basic problem is within the RCBRT for the multivehicle problem; however, if there are multiple pursuers then it may be possible for them to coordinate and herd the evader(s) to force a win, but this strategy would not be captured by pairwise consideration. Consequently, the analytic solution for the basic problem can generate an under approximation for this multivehicle case.

3.4 Unicycle Kinematics and/or Dynamics

The dynamics for the kinematic unicycle model are still given by (1), except that now the evader has two inputs (v_e, ω_e) and the pursuer has two inputs (v_p, ω_p) .

For vehicles with non-trivial inertia the obvious extension beyond a kinematic model is to use linear accelerations as the inputs:

$$\dot{x} = \frac{d}{dt} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix} = \begin{bmatrix} -x_4 + x_5 \cos x_3 + \omega_e x_2 \\ x_5 \sin x_3 - \omega_e x_1 \\ \omega_p - \omega_e \\ a_e \\ a_p \end{bmatrix},$$

where the evader's and pursuer's linear velocities are now part of the state vector (x_4 and x_5 respectively), and the linear accelerations (a_e and a_p) are the inputs (along with ω_e and ω_p). A further extension shifts the angular velocities into the state vector and introduces the angular accelerations as inputs (although this level of model is less often necessary because many robots have comparatively small moments of inertia). It is therefore possible to create versions of RCBRT problems in three, five or seven dimensions in relative coordinates (six, eight or ten dimensions in absolute coordinates) with two inputs for each of the evader and pursuer.

The kinematic and particularly the dynamic unicycle are more realistic models of real world robots, so efficient methods of approximating and representing their RCBRTs would be highly beneficial. Unfortunately, there is no obvious way to extend Merz's results to these higher dimensional cases, so an analytic solution is not available.

4 Approximations from Hamilton-Jacobi-Isaacs

We describe a formulation of RCBRTs as the sublevel sets of the viscosity solution of a Hamilton-Jacobi-Isaacs (HJI) partial differential equation (PDE) in [Mitchell, 2002, Mitchell et al., 2005], and provide open source MATLAB code based on level set methods to approximate these solutions in TOOLBOXLS [Mitchell and Templeton, 2005, Mitchell, 2007]. These approximations may prove useful as another representation of the proposed benchmark RCBRTs against which to compare prospective algorithms, although their memory and computational cost scales exponentially with dimension. Code for the various extension problems will be added as it becomes available.

4.1 Approximating the Basic Problem's RCBRT

We start the section with a very brief sketch of the HJI formulation for the basic problem. A more in-depth presentation of the formulation of this particular problem can be found in [Mitchell, 2002, section 3.1], and implementation details are discussed in [Mitchell, 2007, section 2.6].

The sets are represented by an implicit surface function $\phi : [-T, 0] \times \mathcal{X} \rightarrow \mathbb{R}$ such that

$$\begin{aligned}\mathcal{T} &= \{x \in \mathcal{X} \mid \phi(0, x) \leq 0\}, \\ \mathcal{R}([-t, 0], \mathcal{T}) &= \{x \in \mathcal{X} \mid \phi(t, x) \leq 0\},\end{aligned}$$

for any $t \in [0, T]$. We construct $\phi(0, \cdot)$ from the definition of $\mathcal{T}$; in this case

$$\phi(0, x) = \sqrt{x_1^2 + x_2^2} - \beta. \quad (7)$$

The implicit surface function at times $t > 0$ is the solution of the HJI PDE

$$D_t \phi(t, x) + \min[0, H(x, D_x \phi(t, x))] = 0.$$

In this case the Hamiltonian is given by

$$\begin{aligned}H(x, p) &= - \left(\max_{\omega_e \in [-\bar{\omega}, \bar{\omega}]} \min_{\omega_p \in [-\bar{\omega}, \bar{\omega}]} \begin{bmatrix} p_1 & p_2 & p_3 \end{bmatrix} \begin{bmatrix} -v_e + v_p \cos x_3 + \omega_e x_2 \\ v_p \sin x_3 - \omega_e x_1 \\ \omega_p - \omega_e \end{bmatrix} \right), \\ &= - \left(\max_{\omega_e \in [-\bar{\omega}, \bar{\omega}]} \min_{\omega_p \in [-\bar{\omega}, \bar{\omega}]} \begin{bmatrix} p_1(-v_e + v_p \cos x_3 + \omega_e x_2) \\ +p_2(v_p \sin x_3 - \omega_e x_1) + p_3(\omega_p - \omega_e) \end{bmatrix} \right), \\ &= p_1(v_e - v_p \cos x_3) - p_2(v_p \sin x_3) \\ &\quad - \left(\max_{\omega_e \in [-\bar{\omega}, \bar{\omega}]} \min_{\omega_p \in [-\bar{\omega}, \bar{\omega}]} [\omega_e(p_1 x_2 - p_2 x_1 - p_3) + \omega_p p_3] \right), \\ &= p_1(v_e - v_p \cos x_3) - p_2(v_p \sin x_3) - \bar{\omega} |p_1 x_2 - p_2 x_1 - p_3| + \bar{\omega} |p_3|. \end{aligned} \quad (8)$$

For $v_e = v_p = \bar{\omega} = +1$, this results in the relatively simple $H(x, p) = p_1(1 - \cos x_3) - p_2 \sin x_3 - |p_1 x_2 - p_2 x_1 - p_3| + |p_3|$. There are a number of additional parameters to the level set algorithms which must be chosen before TOOLBOXLS can compute an approximation of this PDE—most notably bounds on the partial derivatives of H with respect to p —but code is provided and the results are illustrated in figure 2.




Figure 7: The pursuer's choice RCBRT for the two dimensional no-heading version. The outer red solid curve is the RCBRT, the inner dotted blue circle is the target set, and the intermediate magenta dashed curves are slices of the analytic solution (the same slices as figure 3, but with x_1 and x_2 swapped).

4.2 Approximating the Extension RCBRTs

In theory the HJI formulation is general enough to represent the RCBRTs from all of the extensions in section 3, but because level set methods require gridding the state space computation is only feasible in state space dimensions beyond four or five by various reformulations, such as those described in [Mitchell and Tomlin, 2003, Mitchell, 2011, Chen et al., 2018, Lee et al., 2019].

The No-Heading Extension. For this extension from section 3.1, $\mathcal{X} = \mathbb{R}^2$ (without the relative heading dimension) and we can easily compute an approximation of the corresponding HJI PDE. The target set is defined as in (7). We start from the Hamiltonian (8) and make the following adjustments:

- Any term involving p_3 is set to zero. The justification for this cancellation can be found in [Mitchell and Tomlin, 2003].
- The value of x_3 in any term is chosen to minimize (to give the advantage to the pursuer) the inner product term. In particular, that means finding

$$\min_{\theta \in [-\pi, \pi)} -v_p(p_1 \cos(\theta) - p_2 \sin(\theta)). \quad (9)$$

A quick bit of calculus can find that the optimal θ^* satisfies $\tan(\theta^*) = -p_2/p_1$, so we evaluate the functional for both possible values of $\theta^* \in [-\pi, \pi)$ and substitute the one which generates the minimum into the Hamiltonian.

The resulting Hamiltonian is

$$H(x, p) = p_1 v_e - v_p [p_1 \cos(\theta^*) - p_2 \sin(\theta^*)] - \bar{\omega} |p_1 x_2 - p_2 x_1|,$$

and figure 7 shows the resulting two dimensional RCBRT. A few things to note about this version of the problem:

- Although it is not obvious in the derivation above, the pursuer's angular velocity ω_p disappears from the problem. Some intuition for this outcome can be gleaned from the fact that we are choosing the relative heading $x_3 \in [-\pi, \pi)$ to be the best value for the pursuer. If the pursuer can always choose the best relative heading, then its turn rate is irrelevant.

- If we seek to give the advantage of the unknown heading to the evader then we can find the θ which maximizes the functional in (9). It is again not obvious, but for the given initial conditions (7) the resulting Hamiltonian is strictly positive, the solution to the resulting HJI PDE is $\phi(t, x) = \phi(0, x)$, and the RCBRT is exactly the target circle in figure 7.

Other Extension Cases. We will add HJI PDE approximations for these cases as they become available.

5 Conclusions

We have described a robust controlled backward reach tube (RCBRT) problem in three dimensions with nonlinear, nonconvex dynamics involving adversarial inputs. The resulting RCBRT is nonconvex and nonsmooth in places. A closed form solution for points on the boundary of this RCBRT is known. From this closed form solution we know that optimal trajectories leading from a boundary point to the target set may be nonunique, may share segments and terminal points with optimal trajectories arising from other points on the boundary, and require only piecewise constant but often bang-bang control signals (a detailed list of the challenging features of this problem can be found in remark 2 in section 2.2). Consequently, this problem appears to be an excellent benchmark for pushing the limits of reachability algorithms.

We listed several extensions to this problem, most of which involve higher state space dimension and some of which have solutions or approximations which can be derived from the analytic solution to the basic problem. Finally, we sketched out an existing algorithm for approximating the solution to these RCBRTs based on HJI PDEs which is quite general but is not scalable. This algorithm is demonstrated on two RCBRT problems, and more can be added depending on interest and computational resources.

Impressive advances have been made in reachability algorithms in the past few years, but most have been limited to systems with a single (or even no) input. We hope that this benchmark proposal will spur further work on robust controlled reachability.

References

- [Aubin et al., 2011] Aubin, J.-P., Bayen, A. M., and Saint-Pierre, P. (2011). *Viability Theory: New Directions*. Systems & Control: Foundations & Applications. Springer.
- [Bak et al., 2019] Bak, S., Tran, H.-D., and Johnson, T. T. (2019). Numerical verification of affine systems with up to a billion dimensions. In *ACM International Conference on Hybrid Systems: Computation and Control (HSCC)*, pages 23–32.
- [Cardaliaguet et al., 1999] Cardaliaguet, P., Quincampoix, M., and Saint-Pierre, P. (1999). Set-valued numerical analysis for optimal control and differential games. In Bardi, M., Raghavan, T. E. S., and Parthasarathy, T., editors, *Stochastic and Differential Games: Theory and Numerical Methods*, volume 4 of *Annals of International Society of Dynamic Games*, pages 177–247. Birkhäuser.
- [Chen et al., 2018] Chen, M., Herbert, S. L., Vashishtha, M. S., Bansal, S., and Tomlin, C. J. (2018). Decomposition of reachable sets and tubes for a class of nonlinear systems. *IEEE Transactions on Automatic Control*, 63(11):3675–3688.
- [Dubins, 1957] Dubins, L. E. (1957). On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents. *American Journal of Mathematics*, 79(3):497–516.

- [Fisac et al., 2018] Fisac, J. F., Akametalu, A. K., Zeilinger, M. N., Kaynama, S., Gillula, J., and Tomlin, C. J. (2018). A general safety framework for learning-based control in uncertain robotic systems. *IEEE Transactions on Automatic Control*, 64(7):2737–2752.
- [Isaacs, 1999] Isaacs, R. (1999). *Differential Games: A Mathematical Theory with Applications to Warfare and Pursuit, Control and Optimization*. Dover. Reprints 1965 edition.
- [Kaynama et al., 2015] Kaynama, S., Mitchell, I. M., Oishi, M. M. K., and Dumont, G. A. (2015). Scalable safety-preserving robust control synthesis for continuous-time linear systems. *IEEE Transactions on Automatic Control*, 60(11):3065–3070.
- [Lee et al., 2019] Lee, D., Chen, M., and Tomlin, C. J. (2019). Removing leaking corners to reduce dimensionality in Hamilton-Jacobi reachability. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 9320–9326.
- [Merz, 1972] Merz, A. W. (1972). The game of two identical cars. *Journal of Optimization Theory and Applications*, 9(5):324–343.
- [Mitchell, 2001] Mitchell, I. (2001). Games of two identical vehicles. Technical Report SUDAAR 740, Department of Aeronautics and Astronautics, Stanford University, Stanford, CA.
- [Mitchell, 2002] Mitchell, I. M. (2002). *Application of Level Set Methods to Control and Reachability Problems in Continuous and Hybrid Systems*. PhD thesis, Scientific Computing and Computational Mathematics Program, Stanford University.
- [Mitchell, 2007] Mitchell, I. M. (2007). A toolbox of level set methods (version 1.1). Technical Report TR-2007-11, Department of Computer Science, University of British Columbia, Vancouver, BC, Canada.
- [Mitchell, 2011] Mitchell, I. M. (2011). Scalable calculation of reach sets and tubes for nonlinear systems with terminal integrators: A mixed implicit explicit formulation. In Frazzoli, E. and Grosu, R., editors, *ACM International Conference on Hybrid Systems: Computation and Control (HSCC)*, pages 103–112. ACM.
- [Mitchell et al., 2005] Mitchell, I. M., Bayen, A. M., and Tomlin, C. J. (2005). A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games. *IEEE Transactions on Automatic Control*, 50(7):947–957.
- [Mitchell and Templeton, 2005] Mitchell, I. M. and Templeton, J. A. (2005). A toolbox of Hamilton-Jacobi solvers for analysis of nondeterministic continuous and hybrid systems. In Morari, M. and Thiele, L., editors, *ACM International Conference on Hybrid Systems: Computation and Control (HSCC)*, number 3414 in Lecture Notes in Computer Science, pages 480–494. Springer Verlag.
- [Mitchell and Tomlin, 2003] Mitchell, I. M. and Tomlin, C. J. (2003). Overapproximating reachable sets by Hamilton-Jacobi projections. *Journal of Scientific Computing*, 19(1–3):323–346.
- [Oishi et al., 2008] Oishi, M., Mitchell, I., Bayen, A. M., and Tomlin, C. J. (2008). Invariance-preserving abstractions: Application to user-interface design. *IEEE Transactions on Control Systems Technology*, 16(2):229–244.
- [Patsko and Turova, 2009] Patsko, V. S. and Turova, V. L. (2009). From Dubins’ car to Reeds’ and Shepp’s mobile robot. *Computing and Visualization in Science*, 12(7):345–364.
- [Reeds and Shepp, 1990] Reeds, J. and Shepp, L. (1990). Optimal paths for a car that goes both forwards and backwards. *Pacific Journal of Mathematics*, 145(2):367–393.
- [Yousefi et al., 2019] Yousefi, M., van Heusden, K., Mitchell, I. M., and Dumont, G. A. (2019). Model-invariant viability kernel approximation. *Systems and Control Letters*, 127:13 – 18.